

Finding Any Nontrivial Coarse Correlated Equilibrium Is Hard

Siddharth Barman*

Katrina Ligett†

Abstract

One of the most appealing aspects of the (coarse) correlated equilibrium concept is that natural dynamics quickly arrive at approximations of such equilibria, even in games with many players. In addition, there exist polynomial-time algorithms that compute exact (coarse) correlated equilibria. In light of these results, a natural question is how *good* are the (coarse) correlated equilibria that can arise from any efficient algorithm or dynamics.

In this paper we address this question, and establish strong negative results. In particular, we show that in multiplayer games that have a succinct representation, it is NP-hard to compute any coarse correlated equilibrium (or approximate coarse correlated equilibrium) with welfare strictly better than the *worst* possible. The focus on succinct games ensures that the underlying complexity question is interesting; many multiplayer games of interest are in fact succinct. Our results imply that, while one can efficiently compute a coarse correlated equilibrium, one cannot provide any nontrivial welfare guarantee for the resulting equilibrium, unless $P = NP$. We show that analogous hardness results hold for correlated equilibria, and persist under the egalitarian objective or Pareto optimality.

To complement the hardness results, we develop an algorithmic framework that identifies settings in which we can efficiently compute an approximate correlated equilibrium with near-optimal welfare. We use this framework to develop an efficient algorithm for computing an approximate correlated equilibrium with near-optimal welfare in aggregative games.

1 Introduction

Equilibria are central solution concepts in game theory, and questions related to the complexity of equilibrium computation have formed a major thread of research in algorithmic game theory. Arguably the most important equilibrium concepts are the Nash equilibrium [19], correlated equilibrium [2], and coarse correlated equilibrium [11]. These solution concepts denote distributions over players' action profiles at which no player can benefit by unilateral deviation, and hence represent stable choices of distributions over player actions. Specifically, a Nash equilibrium is defined to be a product of independent distributions (one for each player); correlated and coarse correlated equilibria are general (joint) probability distributions (see Section 2 for formal definitions).

*California Institute of Technology. E-mail: barman@caltech.edu.

†California Institute of Technology. E-mail: katrina@caltech.edu.

While computation of Nash equilibria has in recent years been shown to be computationally hard, even in games with two players [8], the news for correlated equilibria (CE) and coarse correlated equilibria (CCE) has been more positive. Even in games with many players, there exist a number of natural dynamics that quickly converge to these solution concepts (see, e.g., [17, 10, 12, 6]). In particular, these dynamics induce efficient computation of approximate¹ CE and CCE in multiplayer games; by contrast, computation of approximate Nash equilibria is computationally hard in multiplayer games [22]. In fact, *exact* CE and CCE are efficiently computable in many classes of multiplayer games [21, 14].

Another significant thread of research in algorithmic game theory has been the study of the *quality* of equilibria, often as measured by the social welfare of the equilibrium or its ratio to the social welfare of the socially optimal outcome (c.f. the extensive literature on the price of anarchy (PoA) [20]). Given that we know it is possible to efficiently compute CE and CCE, it is natural to ask *how good* are the equilibria we can efficiently compute? For example, do existing efficient dynamics find the best such equilibria, or at least ones that approximately optimize the social welfare? Since the gap between the worst and the best equilibria (CE or CCE), in terms of social welfare, can be large in natural games (see, e.g., [16, 5]), it is interesting to understand if there exist efficient dynamics or algorithms that avoid—at least to some extent—the bad outcomes. More generally, one can pose the question of efficiently finding CE and CCE *that optimize an objective* (such as the sum of players’ utilities, i.e., the social welfare).

In their notable work, Papadimitriou and Roughgarden [21] show that determining a *socially optimal* CE is NP-hard, in a number of succinct multiplayer games. This result intuitively follows from the fact that determining an action profile with maximum welfare—i.e., solving the problem of welfare optimization even without equilibrium constraints—is NP-hard in general. The hardness result of [21] leaves open the question of computing *near-optimal* CE/CCE, i.e., whether there exist efficient algorithms that compute CE/CCE with welfare at least, say, α times the optimal, for a nontrivial approximation ratio $\alpha \leq 1$. This question forms the basis of the present work.

Technical Aside (succinct games): We note that in general multiplayer games the size of the normal form representation, N , is exponentially large in the number of players; one can compute a CE/CCE that optimizes a linear objective by solving a linear program of size polynomial in N , and hence the computational complexity of equilibrium computation is not interesting for general games. However, most games of interest—such as graphical games, polymatrix games, congestion games, local effect games, network design games, anonymous games, and scheduling games—admit a succinct representation (wherein the above-mentioned linear program can be exponentially large in the size of the representation), and hence it is such succinctly representable games that we (and previous works) study.²

¹A probability distribution over the players’ action profiles is said to be an ε -approximate equilibrium if for any player unilaterally deviating increases utility, in expectation, by at most ε .

²Note that the optimization problem does not become simpler if, instead of a succinct game, one is given access

Results In this paper we establish that, unless $P = NP$, there does not exist any efficient algorithm that computes a CCE with welfare better than the *worst possible* CCE, in succinct multiplayer games (Theorem 1). We also establish similar hardness results for computing equilibria under the egalitarian objective or Pareto-optimality.

Analogous hardness results hold for CE. We note that a classical interpretation of a CE is in terms of a mediator who has access to the players' payoff functions and who draws outcomes from a correlated equilibrium's joint distribution over player actions and privately recommends the corresponding actions to each player. The equilibrium conditions ensure that no player can benefit in expectation by unilaterally deviating from the recommended actions. Therefore, the problem we study here is exactly the computational complexity of the problem that a mediator faces if she wishes to maximize social welfare.

We also extend the hardness result to approximate CE and CCE (Theorem 3). Therefore, while one can efficiently compute an approximate CE/CCE in succinct multiplayer games, one cannot provide any nontrivial welfare guarantees for the resulting equilibrium (unless $P = NP$).

In addition, we show that this hardness result also holds specifically for potential games (generally considered to be a very tractable class of games), and persists even in settings where the gap between the best and worst equilibrium is large.

We note that in these results, the hardness is not simply borrowed from welfare maximization;³ even if the underlying game admits a nontrivial multiplicative approximation for welfare maximization, the problem of determining a CCE with welfare arbitrarily better than the worst CCE remains hard. Another relevant observation is that there always exists an optimal CE/CCE with support size polynomial in the number of players and the number of actions per player.⁴ Therefore, the fact that in multiplayer games there might exist CE/CCE with exponentially large support size does not, in and of itself, account for the complexity of this problem.

We complement these hardness results by developing an algorithmic framework for computing an ε -approximate CE with welfare that is additively ε close to the optimal. This framework establishes a sufficient condition under which the above-mentioned complexity barriers can be circumvented. In particular, we show that if in a given game we can efficiently obtain an *additive* approximation for a *modified-welfare maximization problem*, then we can efficiently compute an approximate CE with high welfare. The modified welfare under consideration can be thought of as a Lagrangian corresponding to the equilibrium constraints (see Definition 7), and the modified-welfare maximization problem entails finding an action profile that maximizes this modified welfare. Note that

to a game via a black box which, when given an action profile a as a query, returns the utilities of all the players at a .

³Welfare maximization refers to the optimization problem of finding an action profile (not necessarily an equilibrium) with maximum possible welfare.

⁴This follows from the fact that CE/CCE are defined by a polynomial number of linear constraints. That is, the set of CE/CCE form a polytope that is defined by a polynomial number of linear inequalities. An optimal CE/CCE is an extreme point of this polytope, and hence its support size is polynomially bounded.

even if welfare (specified by the given utilities) is nonnegative, modified welfare can be negative for certain action profiles. This notably differentiates welfare maximization and modified-welfare maximization, and provides an idea of the technical challenges that one faces when approximating the modified-welfare maximization problem. (Recall that typical multiplicative-approximation techniques cannot handle negative quantities.) Hence, in a given game, the problem of (nontrivially) approximating the modified-welfare maximization problem can be hard, even if the game admits a nontrivial multiplicative-approximation for welfare maximization.

Further, we instantiate this algorithmic framework to compute high-welfare approximate CE in *aggregative games*. These are games wherein the utility of each player is a function of her own action and an aggregate (a constant-dimensional summary vector) of all players' actions; see Section 4.1 for a formal definition. Aggregative games encompass settings like Cournot oligopolies, Bertrand competitions, weighted congestion games, and anonymous games [13, 1, 3, 9]. We develop an efficient additive-approximation algorithm for the modified-welfare maximization problem in aggregative games. Therefore, via the above-mentioned framework, we show how to efficiently compute a high-welfare approximate CE in aggregative games.

Related Work Papadimitriou and Roughgarden [21] showed that the problem of computing an exactly optimal CE is NP-hard for many relevant classes of multiplayer games, including congestion games, graphical games, polymatrix games, local effect games, and scheduling games. Specific instances in which the hardness result of [21] can be completely circumvented, i.e., settings where an *exactly optimal* CE can be efficiently computed, were identified by Jiang and Leyton-Brown [14]. The results in [21] and [14] leave open the question of efficiently computing a CE with *near-optimal* welfare, i.e., the question of *approximating* the optimization problem under consideration. The complexity of this approximation is the focus of our work.

Our main result is negative. In order to prove a positive result for the specific case of computing near-optimal approximate CE in *aggregative games*, we consider a *modified-welfare maximization problem* (MWMP); see Section 4 for a formal definition. Jiang and Leyton Brown [14] consider classes of games in which the MWMP can be solved optimally, and use the ellipsoid method to find an optimal CE. In our setting, exactly solving the MWMP is not computationally feasible (and hence the framework of [14] cannot be applied to aggregative games⁵), but we show that an additive approximation of MWMP suffices to find a near-optimal approximate CE. This entails developing a new algorithm that does not rely on the ellipsoid method.

There is prior work [7, 15, 4] on dynamics that quickly converge to high-welfare CCE in isolated, specific classes of games, such as fair cost sharing games; our results show that it is unlikely that such results can be significantly generalized. Marden et al. [18] develop dynamics that *eventually* converge to Pareto-optimal CCE (see also [23]); these works do not establish polynomial rate of

⁵Knapsack reduces to the problem of welfare maximization in aggregative games.

convergence for the proposed dynamics.

2 Notation

In this paper we consider games with n players and m actions per player. We use A_p to denote the set of actions available to the p th player and A to denote the set of action profiles, $A := \prod_p A_p$. We write $u_p : A \rightarrow [0, 1]$ for the (normalized) utility of player p , and $w : A \rightarrow \mathbb{R}$ is the welfare of an action profile, $w(a) := \sum_{p=1}^n u_p(a)$.⁶ For an action profile $a \in A$, let a_{-p} denote the profile of actions chosen by players other than p . With $A_{-p} := \prod_{q \neq p} A_q$, we have $a_{-p} \in A_{-p}$.

As is typical in the literature, we say that a game is *succinct* if it has an efficient representation. Formally, an n -player m -action game is said to be *succinct* if the player utilities are completely specified via a polynomial-sized string from an input set I . Specifically, for a succinct game, there exists a polynomial (in n and m) time algorithm U that, given a representation $z \in I$ along with a player p and action profile a , returns the utility $u_p(a) = U(z, p, a)$. The game is denoted by $\Gamma(z)$. Many important classes of multi-player games are succinct, e.g., symmetric games, anonymous games, local effect games, congestion games, polymatrix games, graphical games, and network design. This paper is focused on succinct games, since this lets us formally treat settings in which the input, i.e., the utilities in the game, can be efficiently represented. Note that the hardness question becomes moot if we consider the normal form representation of an n -player m -action game as our “input,” since in this case the input itself is exponentially large in n and m . Our hardness results imply the intractability of determining high-welfare CCE in games wherein the underlying utilities are specified through a black box.

We denote the set of probability distributions over a set B by $\Delta(B)$. Given a distribution x over the action profiles A , i.e., $x \in \Delta(A)$, we use $u_p(x)$ for the expected utility of player p under distribution x . Similarly, we write $w(x)$ to denote the expected welfare under x .

Definition 1 (Correlated Equilibrium). *A probability distribution $x \in \Delta(A)$ is said to be a correlated equilibrium if for every player p and every actions $i, j \in A_p$ we have*

$$\sum_{a_{-p} \in A_{-p}} [u_p(j, a_{-p}) - u_p(i, a_{-p})] x(i, a_{-p}) \leq 0,$$

where (i, a_{-p}) denotes an action profile in which player p plays action i and the other players play a_{-p} .

Definition 2 (Coarse Correlated Equilibrium). *A probability distribution $x \in \Delta(A)$ is said to be a coarse correlated equilibrium if for every player p and every action $j \in A_p$ we have*

$$\sum_{a \in A} [u_p(j, a_{-p}) - u_p(a)] x(a) \leq 0,$$

⁶When there are multiple games under consideration within a single proof, we annotate the w to indicate to which game it pertains.

where (j, a_{-p}) denotes an action profile in which player p plays action j and the other players play a_{-p} .

Along these lines, the definition of an approximate correlated equilibrium is as follows:

Definition 3 (ε -Correlated Equilibrium). *A probability distribution $x \in \Delta(A)$ is said to be an ε -correlated equilibrium if for every player p and every actions $i, j \in A_p$ we have*

$$\sum_{a_{-p} \in A_{-p}} [u_p(j, a_{-p}) - u_p(i, a_{-p})]x(i, a_{-p}) \leq \varepsilon.$$

Finally, we define ε -coarse correlated equilibrium.

Definition 4 (ε -Coarse Correlated Equilibrium). *A probability distribution $x \in \Delta(A)$ is said to be an ε -coarse correlated equilibrium if for every player p and every action $i \in A_p$ we have*

$$\sum_{a \in A} [u_p(i, a_{-p}) - u_p(a)]x(a) \leq \varepsilon.$$

3 Hardness Results

In this section we show that, given a succinct game, it is NP-hard to compute a CCE with welfare strictly better than the lowest-welfare CCE. In particular, we develop a reduction that shows that the following decision problem is NP-hard.

Definition 5 (NT). *Let Γ be an n -player m -action game with a succinct representation. NT is defined to be the problem of determining whether Γ admits a coarse correlated equilibrium x such that $w(x) > w(x')$. Here x' denotes the worst CCE of Γ , in terms of social welfare w .*

The hardness of NT implies that, under standard complexity-theoretic assumptions, any non-trivial approximation of the the optimization problem (1) is impossible. Specifically, let x^* denote an optimal CCE of a game (i.e, x^* is an optimal solution of the optimization problem (1)) and x' be a CCE with minimum possible welfare. Write $\beta := w(x')/w(x^*)$, i.e., the ratio of the welfare of the worst CCE to that of the best CCE. In games in which a CCE can be computed efficiently, an efficient β -approximate solution of (1) is direct; we can simply return an arbitrary CCE. The hardness of NT implies that no approximation factor better than β can be achieved in general games. A proof of the NP-hardness of NT is detailed below.

$$\begin{aligned} \max \quad & \sum_{p=1}^n u_p(x) \\ \text{subject to} \quad & x \text{ is a CCE} \end{aligned} \tag{1}$$

Theorem 1. *NT is NP-hard in succinct multiplayer games.*

Proof. We start with a succinct game G from a class of games in which computing a welfare-maximizing action profile is NP-hard. Multiple examples of such classes of games are given in [21]. We reduce the problem of determining an optimal (welfare maximizing) action profile in G to solving NT in a modified succinct game G' . When G is an n -player m -action succinct game, we construct a modified game G' by providing an additional action, b_p , to each player $p \in [n]$. G' is therefore an n -player $(m + 1)$ -action game.

Let A denote the set of action profiles of game G ; similarly, let A' be the action profiles of G' . Let $u_p : A \rightarrow [0, 1]$ and $u'_p : A' \rightarrow [0, 1]$ denote the utility of a player p in G and G' , respectively. Along these lines, let $w(\cdot)$ and $w'(\cdot)$ represent the welfare of action profiles in G and G' , respectively. Note that for every action profile $a' \in A' \setminus A$ there exists at least one player p who is playing the augmented action b_p , i.e., $a'_p = b_p$.

Specifically, we start with the following NP-hard problem: given succinct game G and parameter OPT, determine if there exists an action profile $a \in A$ such that $w(a) \geq \text{OPT}$.⁷ The utilities u_p (and hence also w) are given as succinct input. Using them we define u'_p as follows:

1. For every action profile $a \in A$, $u'_p(a) := w(a)/n$. In other words, on action profiles that belong to the original game we construct an identical-interest game.
2. For every action profile $a' \in A' \setminus A$ such that in a' there is *exactly* one player p who is playing the augmented action b_p (i.e., $a'_p = b_p$ for exactly one player p and $a'_q \neq b_q$ for all $q \neq p$), set $u'_p(a') := \text{OPT}/n$ and $u'_q(a') := 0$ for all $q \neq p$.
3. For action profiles $a' \in A' \setminus A$ in which more than one player is playing the augmented action b_p , we set

$$u'_p(a') = \begin{cases} \varepsilon/n & \text{if } a'_p = b_p \\ 0 & \text{otherwise} \end{cases}$$

Here we select ε to satisfy: $\text{OPT} > \varepsilon \geq \text{OPT}/n$.

Note that G' is a succinct game. Specifically, if game G is succinct then, by definition, we have a polynomial-size specification z for G . In addition, there exists an algorithm U that takes as input z , $p \in [n]$, and $a \in A$, and computes the utility of player p at any action profile a , i.e., $u_p(a)$, in polynomial time. Now to obtain a succinct representation for G' we can use z and U (as a subroutine) and compute utilities u'_p for any player p and action profile a in polynomial time.

Say b denotes the action profile wherein each player is playing the augmented action, $b := (b_1, b_2, \dots, b_n)$. The definition of u'_p implies that $w'(b) := \sum_{p=1}^n u'_p(b) = \sum_{p=1}^n \varepsilon/n = \varepsilon$.

⁷Note that here we are considering an NP-hard *decision* problem and, hence, parameter OPT is part of the input.

We will prove that there exists an action profile $a \in A$ (i.e., an action profile in game G) with $w(a) \geq \text{OPT}$ iff there exists a CCE x in G' that satisfies $w'(x) > w'(b)$. This shows that determining if there exists a CCE x such that $w'(x) > w'(b)$ is NP-hard.

To complete the hardness proof for NT we will show that action profile b is a pure Nash equilibrium (and, therefore, a CCE), and that no other CCE in G' has welfare w' less than b .

Suppose a is an optimal action profile in game G , i.e., $a \in A$ and $w(a) \geq \text{OPT}$. Then a is in fact a pure Nash equilibrium in G' . This follows from the fact that $u'_p(a) = w(a)/n \geq \text{OPT}/n$ (G' is identical interest on $a \in A$); hence (i) for any possible deviation $\hat{a}_p \neq b_p$ for player p we have $u'_p(a) = w(a)/n \geq w(\hat{a}_p, a_{-p})/n = u'_p(\hat{a}_p, a_{-p})$. The first inequality holds since a is an optimal action profile in G ; (ii) for deviation b_p , note that $u'_p(a) \geq \text{OPT}/n = u'_p(b_p, a_{-p})$. Therefore, no player can benefit (increase u'_p) by unilaterally deviating from a , thereby proving that a is a pure Nash equilibrium in G' . Overall, we get that if there exists an action profile $a \in A$ with $w(a) \geq \text{OPT}$ then there exists a CCE x (in particular, an optimal action profile a itself) in G' that satisfies $w'(x) > w'(b)$. Recall that $w'(a) \geq \text{OPT} > \varepsilon = w'(b)$.

It remains to show that if there exists a CCE x such that $w'(x) > w'(b)$ then there exists an action profile $a \in A$ with $w(a) \geq \text{OPT}$. We will consider the set of action profiles in the support of x that are also contained in A , i.e., $\text{Supp}(x) \cap A$. A useful observation is that for all $a' \in A' \setminus A$ the welfare w' satisfies: $w'(a') \leq w'(b)$ (recall, $\varepsilon \geq \text{OPT}/n$). This implies that $\text{Supp}(x) \cap A \neq \emptyset$; otherwise, we would have $w'(x) \leq w'(b)$. Write $\pi > 0$ to denote the probability mass of x on the set $\text{Supp}(x) \cap A$; specifically, $\pi := \sum_{a \in A} x(a)$.

Since x is a CCE, deviating to b_p could not increase any player p 's expected utility:

$$\begin{aligned} w'(x) &= \mathbb{E}_{a \sim x}[u'_p(a)] \\ &\geq \mathbb{E}_{a \sim x}[u'_p(b_p, a_{-p})]. \end{aligned}$$

We can rewrite the above inequality as follows: $\mathbb{E}_{a \sim x}[u'_p(a) - u'_p(b_p, a_{-p})] \geq 0$. Next we expand in terms of conditional expectation

$$\mathbb{E}_{a' \sim x}[u'_p(a') - u'_p(b_p, a'_{-p}) \mid a \in A' \setminus A] \cdot (1 - \pi) + \mathbb{E}_{a \sim x}[u'_p(a) - u'_p(b_p, a_{-p}) \mid a \in A] \cdot \pi \geq 0. \quad (2)$$

Note that for any action profile $a' \in A' \setminus A$ and each player p we have

$$u'_p(a') - u'_p(b_p, a'_{-p}) \leq 0. \quad (3)$$

Either $a'_p = b_p$, in which case $u'_p(a') - u'_p(b_p, a'_{-p}) = 0$; otherwise, $a'_p \neq b_p$ and then $u'_p(a') = 0 < u'_p(b_p, a'_{-p})$.

This implies that the term $\mathbb{E}_{a' \sim x}[u'_p(a') - u'_p(b_p, a'_{-p}) \mid a \in A' \setminus A]$ in inequality (2) is non-positive for every player p . Therefore, the second term in (2), $\mathbb{E}_{a \sim x}[u'_p(a) - u'_p(b_p, a_{-p}) \mid a \in A]$,

must be non-negative for every player p . Summing the second term over all players we get:

$$\mathbb{E}_{a \sim x} \left[\sum_p (u'_p(a) - u'_p(b_p, a_{-p})) \mid a \in A \right] \cdot \pi \geq 0. \quad (4)$$

Recall that $\pi > 0$, i.e., there exists an action profile $a \in A$ such that $x(a) > 0$. Therefore, inequality (4) and the probabilistic method imply that there exists an action profile $a \in A$ such that $\sum_p (u'_p(a) - u'_p(b_p, a_{-p})) \geq 0$. Since $a \in A$, $u'_p(b_p, a_{-p}) = \text{OPT}/n$ for all p . Hence, $\sum_p u'_p(a) \geq \sum_p \text{OPT}/n = \text{OPT}$.

Thus, the existence of a CCE x in G' such that $w'(x) > w'(b)$ implies that there exists an action profile $a \in A$ with $w(a) \geq \text{OPT}$.

To complete the proof, we need to show that b is a pure Nash equilibrium and that no other CCE in G' has welfare w' less than b . The first part of this claim is direct. To prove the second part, suppose by way of contradiction that there existed a CCE x' in G' such that $w'(x') < w'(b)$. Therefore there would exist a player p such that

$$\begin{aligned} \mathbb{E}_{a' \sim x'} [u'_p(a')] &< w'(b)/n \\ &= \varepsilon/n. \end{aligned} \quad (5)$$

But, note that for any $a'_{-p} \in A'_{-p}$ we have $u'_p(b_p, a'_{-p}) \geq \varepsilon/n$. This observation along with inequality (5) implies that p would strictly benefit by unilaterally deviating to b_p . Therefore, x' cannot be a CCE. This completes the proof. \square

Remark: The proof of Theorem 1 can be directly adopted to establish hardness for CE as well. In particular, the fact that any CE x satisfies the inequalities that define a CCE (see Definitions 1 and 2) can be used in the previous proof to show that it is NP-hard to determine a CE with welfare strictly better than the worst possible CE.

In addition, we show below that the reduction given in the proof of Theorem 1 establishes a hardness result for the egalitarian objective as well.

Theorem 2. *In an n -player, m -action succinct game it is NP-hard to determine if there exists a coarse correlated equilibrium x that satisfies $\min_p u'_p(x) > \min_p u'_p(x')$, where u'_p denotes the utility of player p in the given game and x' is the worst equilibrium with respect to the egalitarian objective, i.e., $x' \in \arg \min_{x'' \in \text{CCE}} \{\min_p u'_p(x'')\}$.*

Sketch. Here we use the same notation as in the proof of Theorem 1. Also, as in the previous proof, we obtain a reduction from the following NP-hard problem: given succinct game G , determine if there exists an action profile $a \in A$ such that $w(a) \geq \text{OPT}$.

Note that the action profile b in the constructed game G' is the worst equilibrium with respect to the egalitarian objective, i.e., $b \in \arg \min_{x'' \in \text{CCE}} \{\min_p u'_p(x'')\}$. We can establish this fact by contradiction. In particular, if there existed a CCE x' such that $\min_p u'_p(x') < \min_p u'_p(b) = \varepsilon/n$,

then the player p that obtains the minimum utility under x' could benefit by unilaterally deviating to b_p , contradicting the assumption that x' is a CCE.

To prove this theorem we show that the original game G has an action profile a with $w(a) \geq \text{OPT}$ iff there exists a CCE x such that $\min_p u'_p(x) > \min_p u'_p(b)$. The forward direction follows from the fact that an optimal action profile a with welfare at least OPT is a pure Nash equilibrium in G' . To establish the reverse direction we note that $u'_p(b) = \varepsilon/n$ for all p . Hence if a CCE x satisfies $\min_p u'_p(x) > \min_p u'_p(b)$, then its welfare $w'(x)$ is strictly greater than ε . In other words, $w'(x) > \varepsilon = w'(b)$. But, as shown in the previous proof, this strict inequality suffices to establish the existence of an action profile for which $w(a) \geq \text{OPT}$. Hence, we get the desired claim. \square

Remark: The reduction detailed above also proves that there does not exist a polynomial-time algorithm that computes a Pareto-efficient CCE, unless $\text{P}=\text{NP}$. We can establish this result by noting that a polynomial time algorithm, say \mathcal{A} , that computes any Pareto-efficient CCE can be used to determine whether there exists an action profile a that satisfies $w(a) \geq \text{OPT}$; as before, this suffices to prove the hardness result.

If \mathcal{A} returns b as a Pareto-efficient equilibrium then we know that there does not exist an action profile a such that $w(a) \geq \text{OPT}$, since such an action profile would Pareto dominate b in G' : $u'_p(a) > u'_p(b)$ for all p . Also, note that if \mathcal{A} returns a Pareto-efficient CCE x such that $u'_p(x) = u'_p(b)$ for all p , then again we get that b is Pareto-efficient. So this case is subsumed in the first one. Recall that every CCE x of G' satisfies $u'_p(x) \geq u'_p(b)$. Therefore, the final case entails \mathcal{A} returning a CCE x such that for some p we have $u'_p(x) > u'_p(b)$. Hence, we get that $w'(x) > w'(b)$, which again implies the existence of an action profile a with welfare $w(a) \geq \text{OPT}$.

Remark: Theorems 1 and 2 hold for potential games. This follows from the fact that the reduction used in the proof of these theorems in fact gives us a potential game. Specifically, a potential function ϕ for the constructed game G' is as follows:

1. $\phi(a) := w(a)/n$ for all $a \in A$.
2. For all action profiles $a \in A' \setminus A$ (i.e., in a at least one player is playing its augmented action b_p), we set

$$\phi(a) := \frac{\text{OPT}}{n} + \frac{(k-1)\varepsilon}{n}.$$

Here k is the number of players playing their corresponding augmented action b_p in action profile a , $k = |\{p \mid a_p = b_p\}|$.

A case analysis shows that ϕ is a potential function for G' . In particular, we will show that the following equality holds for each player p and action profiles (a_p, a_{-p}) and (a'_p, a_{-p}) :

$$u'_p(a_p, a_{-p}) - u'_p(a'_p, a_{-p}) = \phi(a_p, a_{-p}) - \phi(a'_p, a_{-p}) \quad (6)$$

Case I: Both (a_p, a_{-p}) and (a'_p, a_{-p}) are action profiles in A . Here we have $u'_p(a_p, a_{-p}) = w(a_p, a_{-p})/n = \phi(a_p, a_{-p})$ and $u'_p(a'_p, a_{-p}) = w(a'_p, a_{-p})/n = \phi(a'_p, a_{-p})$. Hence, in this case (6) holds.

Case II: Action profile $(a_p, a_{-p}) \in A$ and $(a'_p, a_{-p}) \notin A$ (i.e., $a'_p = b_p$). Again, following the definitions of utility u'_p and potential function ϕ we get the equality (6): $u'_p(a_p, a_{-p}) = w(a_p, a_{-p})/n = \phi(a_p, a_{-p})$ along with $u'_p(a'_p, a_{-p}) = \text{OPT}/n = \phi(a'_p, a_{-p})$. The symmetric case of $(a_p, a_{-p}) \notin A$ and $(a'_p, a_{-p}) \in A$ is similarly addressed.

Case III: Both action profiles (a_p, a_{-p}) and (a'_p, a_{-p}) are not in A . If neither a_p nor a'_p is equal to b_p the utility u'_p is zero under both the action profiles. Also, the number of players playing their respective augmented actions b_q is the same in (a_p, a_{-p}) and (a'_p, a_{-p}) , hence $\phi(a_p, a_{-p}) = \phi(a'_p, a_{-p})$. This enforces equality (6).

Now we consider the setting in which exactly one of a_p or a'_p is equal to b_p ; say $a_p = b_p$ (the other possibility (i.e., $a'_p = b_p$) holds by symmetry). Here, $u'_p(a_p, a_{-p}) = \varepsilon/n$ and $u'_p(a'_p, a_{-p}) = 0$. Say $k \in [n]$ is the number of players playing their corresponding augmented action in action profile (a_p, a_{-p}) , then $\phi(a_p, a_{-p}) = \text{OPT}/n + (k-1)\varepsilon/n$ and $\phi(a'_p, a_{-p}) = \text{OPT}/n + (k-2)\varepsilon/n$. Therefore, again, (6) holds.

3.1 Approximate Coarse Correlated Equilibrium

This section establishes the hardness of computing an *approximate* CCE that has high social welfare. Specifically, we consider the problem of computing a $\frac{1}{2n^3}$ -CCE with welfare $(1 + \frac{1}{n})$ times better than the welfare of the worst CCE. Note that there exist regret-based dynamics (c.f [24]) that converge to the set of ε -CCE in time polynomial in $1/\varepsilon$. Therefore, in polynomial time we can compute a $\frac{1}{2n^3}$ -CCE. But, as the following theorem shows, it is unlikely that we can efficiently find a $\frac{1}{2n^3}$ -CCE with any nontrivial welfare guarantee.

Note that in an n -player m -action game a $\frac{1}{2n^3m}$ -CE is guaranteed to be a $\frac{1}{2n^3}$ -CCE (see Definitions 3 and 4). Using this fact, one can directly use the proof given in this section to show that, under standard complexity-theoretic assumptions, there does not exist a polynomial time algorithm that determines a $\frac{1}{2n^3m}$ -CE with any nontrivial welfare guarantee in succinct multiplayer games. It is worth pointing out that in multiplayer games we can always find a $\frac{1}{2n^3m}$ -CE in polynomial time (c.f [24]).

Definition 6 (ANT). *Let Γ be an n -player m -action succinct game. ANT is defined to be the problem of determining whether there exists a $\frac{1}{2n^3}$ -CCE x in Γ such that $w(x) \geq (1 + \frac{1}{n})w(x')$, where x' denotes the worst CCE of Γ , in terms of social welfare w .*

Theorem 3. *In succinct multiplayer games, ANT is NP-hard under randomized reductions: if ANT admits a polynomial-time algorithm then NP admits a polynomial-time randomized algorithm.*

Proof. We will extend the construction presented in the proof of Theorem 1. We start with a game G from a class of games in which it is NP-hard to compute an action profile with welfare within one of the optimal. That is, in G it is NP-hard to compute an action profile a such that $w(a) \geq \max_{a' \in A} w(a') - 1$; note that this is a fairly modest hardness of approximation requirement.

Write $\text{OPT} = \max_{a \in A} w(a)$. Below we develop a polynomial-time randomized algorithm that uses an algorithm for ANT to compute an action profile a that satisfies $w(a) \geq \text{OPT} - 1$. This establishes the stated claim.

To find the desired action profile a , we need a parameter τ that satisfies $\tau \in [\text{OPT} - 1, \text{OPT}]$. Since the utilities in G are normalized between 0 and 1, we have $\text{OPT} \leq n$. Therefore, one of the values in $\{0, 1, \dots, n-1\}$ will give $\tau \in [\text{OPT} - 1, \text{OPT}]$, and we can simply search exhaustively.

Applying the same transformations as in the proof of Theorem 1, we obtain the succinct game G' . While setting utilities in G' we use $\varepsilon = (\tau + 1)/n$, where parameter $\tau \in [\text{OPT} - 1, \text{OPT}]$. Therefore, we have $\frac{\text{OPT}}{n} \leq \varepsilon \leq \frac{\text{OPT}+1}{n}$.

We assume that $\text{OPT} \geq 1$, else finding an action profile a such that $w(a) \geq \text{OPT} - 1$ is trivial. Also, we can assume that $n \geq 4$; recall that for a constant number of players, an optimal CCE can be computed in polynomial time. The following inequality holds under these assumptions: $\text{OPT} \geq (1 + \frac{1}{n}) \frac{\text{OPT}+1}{n}$.

As before, the action profile b is a pure Nash equilibrium, and in fact is a CCE with minimum social welfare.

First, note that an optimal action profile $a^* \in \arg \max_{a \in A} w(a)$ of G is a pure Nash equilibrium (hence, a $\frac{1}{2n^3}$ -CCE) in G' . Also, we have $w'(a^*) = w(a^*) = \text{OPT}$. The bound $w'(a^*) \geq (1 + \frac{1}{n}) w'(b)$ follows from the following chain of inequalities: $\text{OPT} \geq (1 + \frac{1}{n}) \frac{\text{OPT}+1}{n} \geq (1 + \frac{1}{n}) \varepsilon = (1 + \frac{1}{n}) w'(b)$. Thus we get that there exists a $\frac{1}{2n^3}$ -CCE with welfare strictly better than $(1 + 1/n)w'(b)$. This overall ensures that a polynomial-time algorithm for ANT is guaranteed to return a solution. Next we show that any such returned solution can be used to compute an action profile a that satisfies $w(a) \geq \text{OPT} - 1$.

The fact that $w'(a) \leq w'(b)$ for all $a \in A' \setminus A$ and the inequality $w'(x) \geq (1 + \frac{1}{n})w'(b)$ imply that $\sum_{a \in A} w'(a)x(a) \geq \frac{1}{n}w'(b)$. Recall that $w'(b) = \varepsilon \geq \frac{\text{OPT}}{n}$. Therefore, $\sum_{a \in A} w'(a)x(a) \geq \frac{1}{n^2}\text{OPT}$. Since $w(a) = w'(a)$ for all $a \in A$, we have $\max_{a \in A} w'(a) = \text{OPT}$. Therefore, $\pi := \sum_{a \in A} x(a) \geq \frac{1}{n^2}$.

Given that x is a $\frac{1}{2n^3}$ -approximate CCE, analogous to inequality (4) here we have

$$\mathbb{E}_{a \sim x} \left[\sum_p (u'_p(a) - u'_p(b_p, a_{-p})) \mid a \in A \right] \cdot \pi \geq -\frac{1}{2n^2}. \quad (7)$$

Since $\pi \geq \frac{1}{n^2}$, inequality (7) implies $\mathbb{E}_{a \sim x} \left[\sum_p (u'_p(a) - u'_p(b_p, a_{-p})) \mid a \in A \right] \geq -1/2$.

For all $a \in A$ and $p \in [n]$, we have $u'_p(b_p, a_{-p}) = \text{OPT}/n$. Therefore, we get the following bound on the conditional expectation $\mathbb{E}_{a \sim x} \left[\sum_p u'_p(a) \mid a \in A \right] \geq \text{OPT} - 1/2$. For all action profiles $\sum_p u'_p(a) = w'(a) \leq \text{OPT} \leq n$. This implies that in the conditional distribution $\Pr_x(a \mid a \in A)$ the probability mass on action profiles that satisfy $w'(a) \geq \text{OPT} - 1$ is at least $\frac{1}{2n}$.

Therefore, with high probability, we can obtain an action profile that satisfies $w'(a) \geq \text{OPT} - 1$ by drawing polynomially many independent and identically distributed (i.i.d.) samples from the conditional distribution $\Pr_x(a \mid a \in A)$. Since $\pi = \sum_{a \in A} x(a) \geq \frac{1}{n^2}$, we can obtain polynomially many i.i.d. samples from the conditional distribution by drawing polynomially many i.i.d. samples from x . This overall gives us a polynomial-time randomized algorithm to find an action profile that satisfies $w(a) = w'(a) \geq \text{OPT} - 1$. Hence, the stated claim follows. \square

In this section we considered approximate CCE with a specific approximation factor, i.e., we established hardness for $\frac{1}{2n^3}$ -CCE. This was for ease of presentation, and in fact hardness of a parameterized version of ANT can be obtained along the lines of the given proof. In particular, we can show that for any $\delta \in \left[\frac{1}{\text{poly}(n)}, 1 \right]$ it is computationally hard to compute a $\frac{\delta}{2n^2}$ -CCE with welfare greater than $(1 + \delta) w(x')$, where, again, x' denotes the worst CCE.

4 Computing Approximate Correlated Equilibria with Near-Optimal Welfare

In this section, we develop an algorithmic framework for computing an ε -CE with welfare additively ε close to the optimal. The ideas presented in the section can be easily modified to find an ε -CCE with welfare additively ε close to the optimal CCE.⁸

Our framework is based on a novel extension of Blackwell's condition, which is used in the analysis of no-regret algorithms (see, e.g., [24]). The idea here is to define, for each action profile $a \in A$, a vector $r(a)$ whose components list the regret of each player at action profile a . Specifically, for each player $p \in [n]$ and action $j \in [m]$ there is component in $r(a)$ that is equal to $u_p(j, a_{-p}) - u_p(a)$; note that this quantity is the the regret of player p at action profile a with respect to deviation j . The regret vector $r(a)$ has an additional component that is equal to the the difference between the optimal welfare and the welfare of action profile a .

Intuitively, the components of $r(a)$ are defined to ensure that x^* is an optimal CE if and only if the following component-wise inequalities hold: $\mathbb{E}_{a \sim x^*}[r(a)] \leq 0$. Moreover, to find the desired approximate CE it suffices to determine a distribution $x \in \Delta(A)$ that satisfies $\mathbb{E}_{a \sim x}[r(a)] \leq \varepsilon$. Using an extension of Blackwell's condition (see inequality (8)), we develop an algorithm for finding such a distribution x . In particular, via a gradient-descent like argument, we show that action profiles

⁸In order to find an approximate CCE with near-optimal welfare we can define a different regret vector than the one under consideration in this section, whose components are equal to the regret terms that appear in the definition of a CCE. Note that the regret vector for the CCE case is $nm + 1$ dimensional.

satisfying the extended Blackwell condition can be used to determine the desired approximate CE x ; see proof of Theorem 4 for details.

It turns out that finding an action profile that satisfies the extended Blackwell condition corresponds to computing an additive approximation of a modified-welfare maximization problem (see Definition 7). This overall gives us an algorithmic framework that reduces the problem of determining an approximate CE with near optimal welfare to the problem of additive approximating a modified-welfare maximization problem. We instantiate this framework in the context of *aggregative games* in the next section.

Formally, we begin by defining a $d = nm(m-1) + 1$ dimensional regret vector $r(a)$ for each action profile $a \in A$. The first $nm(m-1)$ components of $r(a)$ are indexed by triples (p, i, j) for player $p \in [n]$ and distinct actions $i, j \in [m]$. The (p, i, j) th component of $r(a)$ is equal to $u_p(j, a_{-p}) - u_p(a)$ if $a_p = i$, and is zero otherwise. That is, the (p, i, j) th component is the regret that player p experiences at action profile a by not playing action j . The last (d th) component of $r(a)$ is equal to $w^* - w(a)$. Here w^* denotes the optimal welfare over the set of correlated equilibria, i.e., $w^* := \max\{w(x) \mid x \text{ is a correlated equilibrium}\}$.

Write x^* to denote the welfare-optimal CE, i.e., $w^* = w(x^*)$. Note that for x^* we have that $\mathbb{E}_{a \sim x^*}[r(a)] \leq 0$ holds component-wise. Now, a useful observation is that for any scaling vector $y \in \mathbb{R}_+^d$ with nonnegative components, we have $\mathbb{E}_{a \sim x^*}[y^T r(a)] \leq 0$. Via the probabilistic method, we get that for any $y \in \mathbb{R}_+^d$ there exists an action profile a^* such that

$$y^T r(a^*) \leq 0 \quad (8)$$

Inequality (8) can be thought of as an extension of Blackwell's condition.

This inequality leads us to the objective of maximizing a modified welfare function that is defined as follows.

Definition 7 (Modified Welfare). *Given scaling vector $y \in \mathbb{R}_+^d$ (where the first $nm(m-1)$ components of y are indexed by (p, i, j) for player $p \in [n]$ and actions $i, j \in [m]$ and we refer to the last component of y as y_d), we define modified utilities \tilde{u}_p^y and modified welfare \tilde{w}^y as follows:*

$$\tilde{u}_p^y(a) := y_d u_p(a) + \sum_{j \in A_p} y_{(p, a_p, j)} (u_p(a) - u_p(j, a_{-p})) \quad (9)$$

$$\tilde{w}^y(a) := \sum_p \tilde{u}_p^y(a). \quad (10)$$

For ease of presentation, when y is clear from context we will drop it from the superscript of \tilde{u}_p^y and \tilde{w}^y .

Definition 8 (Modified-Welfare Maximization Problem). *Given a multi-player game and vector $y \in \mathbb{R}_+^d$, the modified-welfare maximization problem (MWMP) is to compute an action profile a of the game that maximizes modified welfare \tilde{w}^y , i.e., the objective is to obtain $\arg \max_{a \in A} \tilde{w}^y(a)$.*

Note that for any vector $y \in \mathbb{R}_+^d$ and any action profile a we have $y^T r(a) = y_d w^* - \tilde{w}^y(a)$. As argued above, for any vector y with non-negative components there exists an action profile a^* that satisfies (8). In particular, a^* satisfies $\tilde{w}^y(a^*) \geq y_d w^*$. Therefore, given a vector y , we can compute an action profile that satisfies (8) by solving an instance of MWMP specified via y . Moreover, an α -additive approximation of MWMP is guaranteed to produce an action profile that satisfies $y^T r(a) \leq \alpha$.

Below we show that (additively) approximating this welfare maximization problem is sufficient to obtain an approximate CE with near-optimal welfare. The hardness result established earlier (see Theorem 3) implies that MWMP cannot be efficiently approximated in general succinct games. However, it is possible for us to approximate MWMP in specific classes of games; in particular, the next subsection details an efficient algorithm to approximate MWMP in *aggregative games*.

Specifically, given a game and vector $y \in \mathbb{R}_+^d$, write $\mathcal{M}(y)$ to denote an $O\left(\frac{\varepsilon^4}{n^4 m^8}\right)$ -additive approximation for MWMP with respect to the specified y . Here, ε is an approximation parameter. Note that an additive approximation $a = \mathcal{M}(y)$ satisfies $y^T r(a) \leq O\left(\frac{\varepsilon^4}{n^4 m^8}\right)$.

Our algorithm, \mathcal{A} , for computing an approximate CE is given below. \mathcal{A} requires access to an additive approximation $\mathcal{M}(y)$ for polynomially many y s. Note that the y s considered during \mathcal{A} 's execution satisfy $y \in [0, n]^d$.

Algorithm for computing ε -correlated equilibrium with near-optimal welfare

Given: an algorithm for computing additive approximation $\mathcal{M}(y)$ in an n -player m -action game;

Return: ε -correlated equilibrium of the game with welfare at least $w^* - \varepsilon$.

- 1: Set a^0 to be an arbitrary action profile of the game and $N = O\left(\frac{n^2 m^4}{\varepsilon^2}\right)$.
 - 2: Let \mathcal{N} denote the negative orthant and $\Pi_{\mathcal{N}}(v)$ denote the Euclidean projection of vector v onto \mathcal{N} .
 - 3: Initialize average regret vector $\bar{r}_0 = r(a^0)$.
 - 4: **for** $t = 1$ to N **do**
 - 5: Set $y = \bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1})$. {Note that the components of y are nonnegative and their magnitude is no more than n }
 - 6: Set $a^t = \mathcal{M}(y)$. {Note that a^t satisfies $y^T r(a^t) \leq O\left(\frac{\varepsilon^4}{n^4 m^8}\right) = O\left(\frac{1}{N^2}\right)$ }
 - 7: Set $\bar{r}_t = \frac{t}{t+1} \bar{r}_{t-1} + \frac{1}{t+1} r(a^t)$.
 - 8: **end for**
 - 9: Return the empirical distribution over the multiset $\{a^0, a^1, a^2, \dots, a^N\}$.
-

Theorem 4. *For a given n -player m -action game, algorithm \mathcal{A} computes an ε -correlated equilibrium with welfare at least $w^* - \varepsilon$. Here w^* denotes the optimal welfare over the set of correlated equilibria of the given game. Moreover, if in the given game additive approximations $\mathcal{M}(y)$ for $y \in [0, n]^d$ can be computed in polynomial (in n , m , and $1/\varepsilon$) time, then \mathcal{A} runs in polynomial time as well.*

Proof. First we establish the stated running-time bound for Algorithm \mathcal{A} . Note that \mathcal{A} iterates $N = O\left(\frac{n^2 m^4}{\varepsilon^2}\right)$ times. Therefore, if additive approximations $\mathcal{M}(y)$ can be computed in polynomial time, then \mathcal{A} runs in polynomial time as well.

Next we establish that \mathcal{A} computes an approximate correlated equilibrium with high welfare. Write x to denote the distribution returned by \mathcal{A} , i.e., x is the empirical distribution over the multi-set of action profiles $\{a^0, a^1, a^2, \dots, a^N\}$. We will show that x satisfies $\mathbb{E}_{a \sim x}[r(a)] \leq \varepsilon$, component-wise. This inequality and the definition of regret vector $r(a)$ imply that x is an ε -correlated equilibrium with welfare at least $w^* - \varepsilon$.

Note that Step (7) of algorithm \mathcal{A} ensures that $\bar{r}_N = \sum_{t=1}^N \frac{1}{N} r(a^t) = \mathbb{E}_{a \sim x}[r(a)]$. Here, the second equality follows from the fact that x is the empirical distribution over action profiles a^1, a^2, \dots, a^N .

Write $d(r, \mathcal{N})$ to denote the Euclidean distance between vector r and the negative orthant \mathcal{N} . The proof proceeds by showing that $d(\bar{r}_N, \mathcal{N})$ is no more than ε . This implies that component-wise \bar{r}_N is no more than ε , and hence we get the desired claim $\mathbb{E}_{a \sim x}[r(a)] \leq \varepsilon$. Recall that \bar{r}_{t-1} denotes the average regret vector considered in the $(t-1)$ th iteration of the algorithm and $\Pi_{\mathcal{N}}(\bar{r}_{t-1})$ denotes the Euclidean projection of this vector onto the negative orthant. The vector $\Pi_{\mathcal{N}}(\bar{r}_{t-1})$ is found by replacing the positive components of \bar{r}_{t-1} by 0, i.e., the i th component of the projection $(\Pi_{\mathcal{N}}(\bar{r}_{t-1}))_i$ is equal to $\min\{0, (\bar{r}_{t-1})_i\}$. We bound the Euclidean distance of \bar{r}_t from the negative orthant as follows:

$$\begin{aligned} d^2(\bar{r}_t, \mathcal{N}) &\leq d^2(\bar{r}_t, \Pi_{\mathcal{N}}(\bar{r}_{t-1})) \\ &= \left\| \frac{t}{t+1} \bar{r}_{t-1} + \frac{1}{t+1} r(a^t) - \Pi_{\mathcal{N}}(\bar{r}_{t-1}) \right\|_2^2 \\ &= \left(\frac{t}{t+1} \right)^2 \|\bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1})\|_2^2 + \left(\frac{1}{t+1} \right)^2 \|r(a^t) - \Pi_{\mathcal{N}}(\bar{r}_{t-1})\|_2^2 \\ &\quad + \frac{2t}{t+1} (\bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1}))^T (r(a^t) - \Pi_{\mathcal{N}}(\bar{r}_{t-1})) \end{aligned} \tag{11}$$

Next we bound the terms on the right-hand side of equality (11). The fact that the utilities of the players are between 0 and 1 implies that for any action profile a the regret vector satisfies $\|r(a)\|_2^2 \leq 2n^2 m^4$. Also, $\|\bar{r}_{t-1}\|_2^2 \leq 2n^2 m^4$, since \bar{r}_{t-1} is an average of regret vectors. Therefore, using the triangle inequality, we get the following bound for the second term in (11), $\left(\frac{1}{t+1} \right)^2 \|r(a^t) - \Pi_{\mathcal{N}}(\bar{r}_{t-1})\|_2^2 \leq \left(\frac{2nm^2}{t+1} \right)^2$.

Step (6) ensures that $(\bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1}))^T r(a^t)$ is no more than $O\left(\frac{1}{N^2}\right)$. In addition, note that the nonzero components of vector $\bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1})$ are the positive components of vector \bar{r}_{t-1} , and on the other hand the nonzero components of vector $\Pi_{\mathcal{N}}(\bar{r}_{t-1})$ are the negative components of \bar{r}_{t-1} . Therefore, we have $(\bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1}))^T \Pi_{\mathcal{N}}(\bar{r}_{t-1}) = 0$. Overall, we get the following bound on the

third term in (11):

$$\begin{aligned} \frac{2t}{t+1} (\bar{r}_{t-1} - \Pi_{\mathcal{N}}(\bar{r}_{t-1}))^T (r(a^t) - \Pi_{\mathcal{N}}(\bar{r}_{t-1})) &\leq O\left(\frac{1}{N^2}\right) \\ &\leq \frac{1}{(t+1)^2} \end{aligned}$$

Using the bounds mentioned above and multiplying equation (11) by $(t+1)^2$ we get

$$(t+1)^2 d^2(\bar{r}_t, \mathcal{N}) \leq t^2 d^2(\bar{r}_{t-1}, \mathcal{N}) + O(n^2 m^4).$$

This leads to a telescoping sum for $1 \leq t \leq N$ that overall gives us

$$N^2 d^2(\bar{r}_N, \mathcal{N}) \leq d^2(\bar{r}_1, \mathcal{N}) + O(n^2 m^4 N). \quad (12)$$

Note that $\|\bar{r}_1\|_2^2 \leq O(n^2 m^4)$, therefore $d^2(\bar{r}_1, \mathcal{N}) \leq O(n^2 m^4)$. Hence, inequality (12) gives $N^2 d^2(\bar{r}_N, \mathcal{N}) \leq O(n^2 m^4 N)$. In other words, $d(\bar{r}_N, \mathcal{N}) \leq O(nm^2/\sqrt{N})$. Given that $N = O\left(\frac{n^2 m^4}{\varepsilon^2}\right)$, we get that the Euclidean distance between \bar{r}_N and the negative orthant is at most ε , i.e., $d(\bar{r}_N, \mathcal{N}) \leq \varepsilon$.

As discussed above, the last inequality implies that $\mathbb{E}_{a \sim x}[r(a)] \leq \varepsilon$, where x is the distribution returned by the algorithm \mathcal{A} . Overall, following the argument outlined above, we get the desired claim. \square

Remark: We can adapt this algorithmic framework to the egalitarian objective or Pareto efficiency, instead of welfare maximization.

For the egalitarian objective, for each action profile, instead of regret vector $r(a)$, we can consider $(nm(m-1) + n)$ -dimensional vector $\rho(a)$. The first $nm(m-1)$ components of $\rho(a)$ and $r(a)$ are the same. But, the last n components of $\rho(a)$ are set equal to $w' - u_p(a)$ for each $p \in [n]$. Here w' is the optimal value of the egalitarian objective, $w' := \arg \max_{x \in \text{CE}} \min_p u_p(x)$. Working with $\rho(a)$ and the corresponding modified-welfare function, we can obtain an ε -CE x that satisfies $\min_p u_p(x) \geq w' - \varepsilon$.

To find an approximate correlated equilibrium that is nearly Pareto efficient, we pick a specific player q and replace the last component of the regret vector $r(a)$ by $w'' - u_q(a)$. Here $w'' := \arg \max_{x \in \text{CE}} u_q(x)$. In this case, we can consider the relevant modified-welfare function and overall obtain an ε -CE that satisfies $u_q(x) \geq w'' - \varepsilon$. Since there does not exist a CE wherein the utility of q is ε more than $u_q(x)$, we get that x is ε -Pareto efficient.

4.1 Aggregative Games

This section presents a polynomial-time additive-approximation algorithm for MWMP in aggregative games. An n -player m -action aggregative game with action profiles A is specified by an aggregator function $S : A \rightarrow [-W, W]^k$ and utility-defining functions $v_p : A_p \times [-W, W]^k \rightarrow [0, 1]$.

The function S serves as a sufficient statistic for the utilities of the player; specifically, the utility of player p at action profile a (i.e., $u_p(a)$) is equal to $v_p(a_p, S(a))$. Note that here the utility depends on the action of the player, a_p , and the aggregated vector, $S(a)$. In aggregative games, the function S is additively separable; in particular, there exist vectors $f_p(a_p) \in [-W', W']^k$ for each player $p \in [n]$ and action $a_p \in A_p$ such that for any action profile $a \in A$ we have the following component-wise equality: $S(a) = \sum_p f_p(a_p)$. Here, the dimension k is assumed to be a fixed constant and W and W' are polynomially bounded in n and m .

Along the lines of prior work (see [3, 9]), we consider the setting in which the influence of the aggregator on the utilities is bounded: $|v_p(a_p, s) - v_p(a_p, s')| \leq \|s - s'\|_\infty$ for all players $p \in [n]$, actions $a_p \in A_p$, and vectors $s, s' \in [-W, W]^k$. The assumption that k is a fixed constant can be mute without this bounded influence property.

Recall that modified utilities $\tilde{u}_p^y(a)$ are defined in terms of the utilities $u_p(a)$; see Definition 7. In this section we will only consider vectors y whose components are linearly bounded, i.e., $y \in [0, n]^d$; in order to apply Theorem 4 it suffices to consider such linearly-bounded vectors.

We begin by discretizing the aggregating vectors $f_p(a_p)$ such that their components are multiples of parameter δ , which will be set appropriately. That is, for all $p \in [n]$ and $a_p \in A_p$, the components of vector $f_p(a_p)$ are rounded to the nearest multiple of δ . Note that component-wise the vectors $f_p(a_p)$ are polynomially bounded; hence, a polynomially small δ ensures that even after discretization for all action profiles a , the aggregated value $S(a)$ remains within $O\left(\frac{\varepsilon^4}{\text{poly}(n, m)}\right)$ —under the ℓ_∞ norm—of the original (undiscretized) value. The bounded influence assumption, $|v_p(a_p, s) - v_p(a_p, s')| \leq \|s - s'\|_\infty$ and the fact that $y \in [0, n]^d$ ensures that the discretization process does not change the modified utility $\tilde{w}^y(a)$ by more than $O\left(\frac{\varepsilon^4}{\text{poly}(n, m)}\right)$, for any action profile a . Overall, this implies that (with a polynomially small δ) if we compute an action profile a' that maximizes \tilde{w}^y with the discretized aggregator function then a' is an $O\left(\frac{\varepsilon^4}{n^4 m^8}\right)$ -additive approximation for MWMP with the original (undiscretized) aggregator function.

Throughout the remainder of the section we will work with the discretized aggregator. Now, all the discretized vectors $f_p(a_p)$ s are contained in $\left\{0, \pm\delta, \pm2\delta, \pm3\delta, \dots, \pm\left\lceil\frac{W'}{\delta}\right\rceil\delta\right\}^k$. Write \mathcal{G} to denote the k -dimensional grid defined as follows: $\mathcal{G} := \{\sum_{q=1}^p f_q(a_q) \mid \text{for all } p \in [n] \text{ and each } a_q \in A_q\}$. Since δ is polynomially small we have $|\mathcal{G}| = \left(\frac{nm}{\varepsilon}\right)^{O(k)}$. Also, for all action profiles $a \in A$, the discretized aggregator function value $S(a)$ is contained in \mathcal{G} .

We develop a dynamic program that works over \mathcal{G} and computes an action profile that maximizes the modified welfare \tilde{w}^y . As discussed above, this gives us an $O\left(\frac{\varepsilon^4}{n^4 m^8}\right)$ -additive approximation for MWMP.

Our main result of this section is that an additive approximation for MWMP can be computed efficiently when the scaling vector y is contained in $[0, n]^d$.

Theorem 5. *Given an n -player m -action aggregative game and a scaling vector $y \in [0, n]^{(nm(m-1)+1)}$, there exists a polynomial-time algorithm that computes an $O\left(\frac{\varepsilon^4}{n^4 m^8}\right)$ -additive approximation for the*

MWMP instance specified via y .

Proof. Throughout the proof we work with the modified welfare function \tilde{w}^y that is specified by the discretized aggregator function. In particular, via a dynamic program we will compute an action profile that maximizes \tilde{w}^y . As mentioned above, this gives the desired additive-approximation guarantee.

For each vector $s \in \mathcal{G}$ we will maximize the modified welfare \tilde{w}^y function over the set of action profiles $A(s) := \{a \in A \mid S(a) = s\}$ in polynomial time. Write a^* to denote an action profile that maximizes \tilde{w}^y . Discretization ensures that $S(a^*) \in \mathcal{G}$. Also, the fact that the cardinality of \mathcal{G} is polynomially bounded implies that efficiently optimizing over $A(s)$ for each $s \in \mathcal{G}$ gives an action profile that maximizes \tilde{w}^y in polynomial time.

The remainder of the proof details an algorithm that, given vector $s \in \mathcal{G}$, solves the following optimization problem in polynomial time: $\arg \max_{a \in A(s)} \tilde{w}^y(a)$.

First, we define modified utility function $\tilde{v}_p(a_p, s)$ in terms of the given functions $v_p(a_p, s)$. Recall that the vector y is $d = nm(m-1) + 1$ dimensional and its first $nm(m-1)$ components are indexed by triples (p, i, j) with $p \in [n]$ and distinct $i, j \in [m]$. Using vector y as a parameter we define,

$$\tilde{v}_p^y(a_p, s) := y_d v_p(a_p, s) + \sum_{j \in A_p} y_{(p, a_p, j)} \cdot [v_p(a_p, s) - v_p(j, s - f_p(a_p) + f_p(j))]$$

A key observation is that for any action profile a , if $s = S(a)$ then $\tilde{u}_p^y(a) = \tilde{v}_p^y(a_p, s)$ and, hence, $\tilde{w}^y(a) = \sum_{p=1}^n \tilde{v}_p^y(a_p, s)$. Therefore, by the definition of $A(s)$, for all $a \in A(s)$ the following equality holds $\tilde{w}^y(a) = \sum_{p=1}^n \tilde{v}_p^y(a_p, s)$. Hence, $\arg \max_{a \in A(s)} \sum_{p=1}^n \tilde{v}_p^y(a_p, s)$ is equal to $\arg \max_{a \in A(s)} \tilde{w}^y(a)$.

We solve $\arg \max_{a \in A(s)} \sum_{p=1}^n \tilde{v}_p^y(a_p, s)$ via a dynamic program that fills a matrix $M(p, s')$ indexed by $p \in [n]$ and $s' \in \mathcal{G}$. In particular, $M(p, s')$ is set equal to $\max_{a_1, a_2, \dots, a_p} \{\sum_{q=1}^p \tilde{v}_q^y(a_q, s) \mid \sum_{q=1}^p f_q(a_q) = s'\}$. Here, the entry $M(n, s)$ is equal to the target optimal value $\max_{a \in A(s)} \tilde{w}^y(a)$. We can initialize $M(1, s')$ by going over actions in A_1 ; specifically, $M(1, s') = \max_{a_1} \{\tilde{v}_1^y(a_1, s) \mid f_1(a_1) = s'\}$. In general, we use the recurrence relation $M(p, s') = \max_{a_p, s''} \{\tilde{v}_p^y(a_p, s) + M(p-1, s'') \mid f_p(a_p) + s'' = s'\}$ to complete the matrix. A direct inductive argument proves the correctness of this dynamic program.

Since, $|\mathcal{G}|$ is polynomially bounded, the size of the matrix is also polynomially bounded. Overall, the dynamic program runs in polynomial time and the stated claim follows. \square

Theorem 5 shows that the additive approximation required in Theorem 4 can be computed in polynomial time. Therefore, the two theorems together imply that in aggregative games an approximate correlated equilibrium with near-optimal welfare can be computed in polynomial time.

Corollary 1. *In n -player m -action aggregative games an ε -correlated equilibrium with welfare $w^* - \varepsilon$ can be computed in time polynomial in n , m , and $1/\varepsilon$. Here w^* denotes the optimal welfare over the set of correlated equilibria of the given game.*

Acknowledgements

The authors thank Nikhil Bansal for helpful discussions. This work was supported by NSF grants CNS-0846025, CCF-1101470, and CNS-1254169, along with a Microsoft research faculty fellowship, a Google faculty research award, and a Linde/SISL postdoctoral fellowship. Katrina Ligett gratefully acknowledges the support of the Charles Lee Powell Foundation.

References

- [1] Daron Acemoglu and Martin Kaae Jensen. Aggregate comparative statics. *Games and Economic Behavior*, 81:27–49, 2013.
- [2] Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96, 1974.
- [3] Yakov Babichenko. Best-reply dynamic in large aggregative games. *Available at SSRN 2210080*, 2013.
- [4] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Circumventing the price of anarchy: Leading dynamics to good behavior. *SIAM Journal on Computing*, 42(1):230–264, 2013.
- [5] Vittorio Bilo, Michele Flammini, and Luca Moscardelli. The price of stability for undirected broadcast network design with fair cost allocation is constant. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 638–647. IEEE, 2013.
- [6] Avrim Blum and Yishay Mansour. From external to internal regret. *The Journal of Machine Learning Research*, 8:1307–1324, 2007.
- [7] Moses Charikar, Howard Karloff, Claire Mathieu, Joseph Seffi Naor, and Michael Saks. Online multicast with egalitarian cost sharing. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures*, pages 70–76. ACM, 2008.
- [8] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):14, 2009.
- [9] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. Privacy and truthful equilibrium selection for aggregative games. *arXiv preprint arXiv:1407.7740*, 2014.
- [10] Dean P Foster and Rakesh V Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- [11] James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.

- [12] Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.
- [13] Martin Kaae Jensen. Aggregative games and best-reply potentials. *Economic Theory*, 43(1):45–66, 2010.
- [14] Albert Xin Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior*, 2013.
- [15] Robert Kleinberg, Georgios Piliouras, and Éva Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 533–542. ACM, 2009.
- [16] Euiwoong Lee and Katrina Ligett. Improved bounds on the price of stability in network cost sharing games. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 607–620. ACM, 2013.
- [17] Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.
- [18] Jason R Marden, H Peyton Young, and Lucy Y Pao. Achieving pareto optimality through distributed learning. In *CDC*, pages 7419–7424, 2012.
- [19] John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.
- [20] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Cambridge University Press, 2007.
- [21] Christos H Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):14, 2008.
- [22] Aviad Rubinstein. Inapproximability of nash equilibrium. *Proceedings of the forty-seventh annual ACM Symposium on Theory of Computing (STOC)*, 2015.
- [23] Yang Song, Starsky HY Wong, and Kang-Won Lee. Optimal gateway selection in multi-domain wireless networks: a potential game perspective. In *Proceedings of the 17th annual international conference on Mobile computing and networking*, pages 325–336. ACM, 2011.
- [24] H Peyton Young. *Strategic learning and its limits*. Oxford University Press, 2004.